arXiv:1407.6665v1 [cs.DS] 27 Jun 2014

A Tight Lower Bound for Decrease-Key in the Pure Heap Model

John Iacono* Özgür Özkan

New York University
Brooklyn, New York, USA

Abstract

We improve the lower bound on the amortized cost of the decrease-key operation in the pure heap model and show that any pure-heap-model heap (that has a $\mathcal{O}(\log n)$ amortized-time extract-min operation) must spend $\Omega(\log \log n)$ amortized time on the decrease-key operation. Our result shows that sort heaps as well as pure-heap variants of numerous other heaps have asymptotically optimal decrease-key operations in the pure heap model. In addition, our improved lower bound matches the lower bound of Fredman [J. ACM 46(4):473-501 (1999)] for pairing heaps [M.L. Fredman, R. Sedgewick, D.D. Sleator, and R.E. Tarjan. Algorithmica 1(1):111-129 (1986)] and surpasses it for pure-heap variants of numerous other heaps with augmented data such as pointer rank-pairing heaps.

1 Introduction

A *priority queue* is a fundamental data structure that supports the following operations to maintain a totally ordered set[1] $\mathcal{S}$: $p =$ Insert(x): inserts the key x into $\mathcal{S}$ and returns a pointer p used to perform future Decrease-Key operations on this key, $x =$ Extract-Min(): removes the minimum item in $\mathcal{S}$ from $\mathcal{S}$ and returns it, and Decrease-Key($p, \Delta x$): reduces the key value pointed to by p by some non-negative amount Δx.

The fast execution of Decrease-Key is vital to the runtime of several algorithms, most notably Dijkstra's algorithm [2]. The constant-amortized-time implementation of the Decrease-Key operation is the defining feature of the *Fibonacci heap* [9].

However, Fibonacci heaps are not pointer model structures in the standard textbook presentation [1] for the following reason: each node is augmented with a $\log \log n$-bit integer (in the range $1 \ldots \log n$), and in the implementation of Extract-Min it is required to separate a number of nodes, call it k, into groups of nodes with like number. This is done classically using bucket sort in time $\mathcal{O}(k + \log n)$. Bucket sort using indirect addressing to access each bucket is distinctly not allowed in the pointer model; realizing this, there is a note in the original paper [9] showing how by adding another pointer from every node to a node representing its bucket, bucket sort can be simulated at no additional asymptotic cost. However, these pointers are non-heap pointers to nodes

*Research supported by NSF Grant CCF-1018370.

[1]The key values can be in any form so long as they are totally ordered and an $O(1)$-time comparison function is provided. This is more permissive than saying they they are comparison-based. We do not restrict algorithms from doing things like making decisions based on individual bits of a key. Some priority queues, including Fibonacci heaps and our sort heaps, also efficiently support the Meld operation where two priority queues are combined into one.

with huge in-degree which store no keys and thus violate the tree requirement of the pure heap model, as well as the spirit of what we usually think of as a heap.

There have been several alternatives to Fibonacci heaps presented with the same amortized runtimes which are claimed to be simpler than the original [14, 5, 10].

All of these heaps also use indirect addressing or non-heap pointers. The rank-pairing heap [10] has the cleanest implementation of all of them, implementing DECREASE-KEY by just disconnecting the node from it parents and not employing anything more complicated like the cascading cuts of Fibonacci heaps. We can easily modify Fibonacci heaps and the aforementioned alternatives to only use heap pointers by simply using $\mathcal{O}(\log \log n)$ time to determine which of the $\mathcal{O}(\log n)$ buckets each of the k keys are in. We call such a variant of rank-pairing heaps a *pointer rank-pairing* heap. However, this alteration increases the time of DECREASE-KEY in Fibonacci heaps and their variants to $\mathcal{O}(\log \log n)$.

In [8], a new heap called the *pairing heap* was introduced. The pairing heap is a self-adjusting heap, whose design and basic analysis closely follows that of splay trees [16]. They are much simpler in design than Fibonacci heaps, and they perform well in practice. It was conjectured at the time of their original presentation that pairing heaps had the same $\mathcal{O}(\log n)$ amortized time EXTRACT-MIN, and $\mathcal{O}(1)$ amortized time INSERT and DECREASE-KEY as Fibonacci heaps, however, at their inception only a $\mathcal{O}(\log n)$ amortized bound was proven for all three operations. Iacono has shown a working-set like runtime bound [11], and also that INSERT does in fact have $\mathcal{O}(1)$ amortized time [12, 11]. However in [7], Fredman refuted the conjectured constant-amortized-time DECREASE-KEY in pairing heaps by proving that pairing heaps have a lower bound of $\Omega(\log \log n)$ on the DECREASE-KEY operation. The result he proved was actually more general: he created a model of heaps which includes both pairing heaps and Fibonacci heaps, and produced a tradeoff between the number of bits of data augmented and a lower bound on the runtime of DECREASE-KEY. Pairing heaps have no augmented bits, and were shown to have a $\Omega(\log \log n)$ amortized lower bound on DECREASE-KEY while in his model a $\mathcal{O}(1)$ DECREASE-KEY requires $\Omega(\log \log n)$ bits of augmented information per node, which is the number of bits of augmented information used by Fibonacci heaps and variants (a number called rank, which is an integer with logarithmic range is stored in each node). More recently, Pettie has shown a upper bound of $\mathcal{O}(2^{2\sqrt{\log \log n}})$ for DECREASE-KEY in pairing heaps [15]. It remains open where in $\Omega(\log \log n) \ldots \mathcal{O}(2^{2\sqrt{\log \log n}})$ the cost of DECREASE-KEY in a pairing heap lies. Elmasry has shown that a simple variant of pairing heaps has $\mathcal{O}(\log \log n)$ amortized time DECREASE-KEY operation [4, 3]. However, this variant is not in Fredman's model and thus the $\Omega(\log \log n)$ lower bound does not apply.

Recently, a new model of pointer model heaps called the *pure heap* model was introduced [13], along with a variant of Elmasry's heap in the model called the sort heap which supports the DECREASE-KEY operation in $\mathcal{O}(\log \log n)$ amortized time. A pure heap is a pointer-based forest of rooted trees, each node holding one key, that obeys the heap property (the key of the source of every pointer is smaller than the key of the destination); the nodes may be augmented, and all operations must be valid in the pointer model; heap pointers are only removed if one of the nodes is removed or if a DECREASE-KEY is performed on the node that a heap pointer points to. This model is both simple and captures the spirit of many heaps like the pairing heap, and is meant to be a clean definition analogous to that of the well-established binary search tree (BST) model [19].

It was shown in [13] that any heap in the pure heap model that has $\mathcal{O}(\log n)$ amortized-time EXTRACT-MIN and INSERT operations must spend $\Omega\left(\frac{\log \log n}{\log \log \log n}\right)$ amortized time on the DECREASE-KEY operation, *no matter how many bits of data each node is augmented with.* We

improve this result by showing that any heap in the pure heap model that has $\mathcal{O}(\log n)$ amortized-time EXTRACT-MIN and INSERT operations must in fact spend $\Omega(\log \log n)$ amortized time on the DECREASE-KEY operation. This bound is asymptotically tight as it matches the upper bounds for numerous heaps including pointer rank-pairing heaps and variants, and sort heaps. It also matches (or surpasses in the case of pointer rank-pairing heaps) the lower bounds achieved in Fredman's model. In this view, Fibonacci heaps are a typical RAM-model structure that squeeze out a $\log \log n$ factor over the best structure in a natural pointer-based model by beating the sorting bound using bit tricks (of which bucket sort is a very primitive example). Thus, Fibonacci heaps have fast DECREASE-KEY is not (only) because of the augmented bits as Fredman's result suggests, but rather because they depart from the pure heap model.

We review the pure heap model in Section 2 and present the improved lower bound in Section 3.

2 The Pure Heap Model

Here we define the *pure heap* model, and how priority queue operations on data structures in this model are executed in it. The pure heap model requires that at the end of every operation, the data structure is an ordered forest of general heaps. Each node is associated with a key $x \in \mathcal{S}$. We will use x both to refer to a key and the node in the heap containing the key. Inside each node is stored the key value, pointers to the parent, leftmost child and right sibling of the node, along with other possible augmented information. The *structure* of the heap is the shape of the forest, without regard to the contents of the nodes. The *position* of a node is its position in the forest of heaps relative to the right. The idea is that location is invariant under adding new siblings to the left.

An algorithm in the pure heap model implements the priority queue operations as follows: INSERT operations are executed at unit cost by adding the new item as a new leftmost heap. DECREASE-KEY is executed at unit cost by disconnecting from its parent the node containing the key to be decreased (if it is not a root), decreasing the key, and then placing it as the leftmost root in the forest of heaps. An EXTRACT-MIN operation is performed by first executing a sequence of pointer-based suboperations which are fully described below. After executing the suboperations, the forest is required to be monoarboral (i.e. have only one heap). Thus, the root of this single tree has as its key the minimum key in $\mathcal{S}$. This node is then removed, the key value is returned, and its children become the new roots of the forest. The cost is the number of suboperations performed.

Note that some data structures are not presented exactly in the framework as described above but can be easily put into this mode by being lazy. For example, in a pairing heap, the normal presentations of INSERT and DECREASE-KEY cause an immediate pairing with the single existing root. However, such pairings can easily be deferred until the next EXTRACT-MIN, thus putting the resulting structure in our pure heap framework. For a more elaborate example of transforming a heap into one based on pairings, see [6].

2.1 Executing an EXTRACT-MIN

To execute an EXTRACT-MIN, the minimum must be determined. In an EXTRACT-MIN, the forest of heaps must be combined into a single heap using an operation called *pairing*, which takes two roots and attaches the root with larger key value as the leftmost child of the root with smaller key value. (Note that while the pairing operation brings to mind pairing heaps, it is the fundamental

building block of many heaps. Even the *skew heap* [17], which seems at first glance to not use anything that looks like the pairing operation, can be shown in all instances to be able to be transformed into a pairing-based structure at no decrease in cost [6].) In the execution of the EXTRACT-MIN operation, the use of some constant number ρ of pointers $p_1, p_2, \ldots p_\rho$ is allowed. They are all initially set to the leftmost root.

1. HASPARENT(i): Return if the node pointed to by p_i has a parent.
2. HASLEFTSIBLING(i): Return if the node pointed to by p_i has a left sibling.
3. HASRIGHTSIBLING(i): Return if the node pointed to by p_i has a right sibling.
4. HASCHILDREN(i): Return if the node pointed to by p_i has any children.
5. COMPARE(i, j): Return if the key value in the node pointed to by p_i is less than or equal to the key value in the node pointed to by p_j.
6. PAIR(i, j): Perform a *pairing* on two pointers p_i and p_j where the tree that p_j points to is attached to p_i as its leftmost subtree. It is a required precondition of this suboperation that both p_i and p_j point to roots, and that this was verified by the HASPARENT suboperation, and that the key value in the node pointed to by p_i is smaller than the key value in the node pointed to by p_j, and that is was verified by the COMPARE suboperation.
7. UNPAIR(i): Disconnect the left child of the node p_i points to and make it the left sibling of p_i. To use this suboperation, it is required that p_i has children, and crucially that p_i and its left child were paired with a PAIR($\cdot, \cdot$) suboperation performed in the same EXTRACT-MIN as this one. Without this restriction, proof of the Lemma 6 below would not hold.
8. SET(i, j): Set pointer p_i to point to the same node as p_j.
9. SWAP(i, j): Swap the heaps pointed to by p_i and p_j.
10. MOVETOPARENT(i): Move a pointer p_i to the parent of the node currently pointed to. It is a precondition of this operation that the node that p_i points to has a parent, and that this was verified by the HASPARENT operation.
11. MOVETOLEFTMOSTCHILD(i): Move the pointer p_i to the leftmost child. It is a precondition of this operation that the node that p_i points to has children, and that this was verified by the HASCHILDREN operation.
12. MOVETORIGHTSIBLING(i): Move the pointer p_i to the sibling to the right. It is a precondition of this operation that the node that p_i points to has a right sibling, and that this was verified by the HASRIGHTSIBLING operation.
13. MOVETOLEFTSIBLING(i): Move a pointer p_i to the sibling to the left It is a required precondition of this operation that the node that p_i points to has a left sibling, and that this was verified by the HASLEFTSIBLING operation.
14. END(): Marks the end of the suboperation sequence for a particular EXTRACT-MIN. It is a required precondition of this operation that the forest of heaps contains only one heap, and that this was verified through the use of the HASPARENT(i), HASLEFTSIBLING(i), and HASRIGHTSIBLING(i) on a pointer p_i that points to the unique root.

Operations 1-4 return a boolean; the remainder have no return value.

The total number of suboperations, including the parameters, is defined to be η. Observe that $\eta = \Theta(\rho^2)$, which is $\Theta(1)$ since ρ is a constant. A sequence of suboperations is a *valid* implementation of the EXTRACT-MIN operation if all the preconditions of each suboperation are met and the last suboperation is an END(). In the pure heap model, the only thing that differentiates between different algorithms is in the choice of the suboperations to execute EXTRACT-MIN operations. In these operations it is the role of the particular *heap algorithm* to specify which suboperations should be performed for each EXTRACT-MIN. There are no restrictions as to how an algorithm determines the suboperation sequence for each EXTRACT-MIN other than the suboperation sequence must be valid; the algorithm need not restrict its actions to information in the nodes visited in that operation.

3 Lower Bound

Theorem 1 (Main Theorem)**.** *In the pure heap model with a constant number of pointers, if* EXTRACT-MIN *and* INSERT *have an amortized cost of* $\mathcal{O}(\log n)$*, then* DECREASE-KEY *has an amortized cost of* $\Omega(\log \log n)$*.*

The proof will follow by contradiction. Assume that there is a pure heap model algorithm $\mathcal{A}$ where EXTRACT-MIN and INSERT have an amortized cost of at most $\alpha \log n$, for some constant α, and DECREASE-KEY has an amortized cost of at most $dc(n)$, for some $dc(n) = o(\log \log n)$. The existence of the algorithm $\mathcal{A}$, the constant α and the function $dc(n)$ will be assumed in the definitions and lemmas that follow. A sufficiently large n is also assumed.

3.1 Preliminaries

The proof is at its core an adversary argument. Such arguments look at what the algorithm has done and then decide what operations to do next in order to guarantee a high runtime. But, our argument is not straightforward as it works on sets of sequences of operations rather than a single operation sequence. There is a hierarchy of things we manipulate in our argument:

Suboperation. The suboperations of §2.1 are the very basic unit-cost primitives that can be used to implement EXTRACT-MIN, the only operation that does not have constant actual cost. It is at this level that definitions have been made to enforce pointer model limitations.

Operation. We use *operation* to refer to a priority queue operation. In this proof, the adversary will only use INSERT, DECREASE-KEY, and EXTRACT-MIN.

Sequence. Operations are combined to form sequences of operations.

Set of operation sequences. Our adversary does not just work with a single operation sequence but rather with sets of operation sequences. These sets are defined to have certain invariants on the heaps that result from running the sequences which bound the size of the sets of operation sequences under consideration.

Evolution. We use the word *evolution* to refer to a function the adversary uses to take a set of operation sequences, and modify it. The modifications performed are to append operations

to sequences, remove sequences, and to create more sequences by taking a single sequence and appending different operations to the end.

Rounds. Our evolutions are structured into *rounds*.

The proof will start with a set containing a single operation sequence, and then perform rounds of evolutions on this set; the exact choice of evolutions to perform will depend on how the algorithm executes the sequences of operations in the set. The evolutions in a round are structured in such a way that most rounds increase the size of the set of operations. After sufficiently many rounds, an upper bound on the maximum size of the set of operation sequences will be exceeded, thus giving a contradiction.

In the remainder of Section 3.1, we provide numerous definitions and present a number of notions needed for the proof. In Section 3.2, we will give a more detailed outline of the proof. In Sections 3.3 and 3.4, we describe our evolutions and rounds as defined above, respectively. Finally, in Sections 3.5 and 3.6 we derive a contradiction to prove our result.

3.1.1 Rank of a node

As in many previous works on heaps and trees, the notion of the *rank* of a node in the heap is vital. The rank of a node is meant to be a rough proxy for the logarithm of the size of the subtree of the node.

While the basic analysis of pairing heaps and splay trees [8, 16] use exactly this as the rank, the definition of rank here is more delicate and is an extension of that used in [7]. As in [7], rank here is always a nonnegative integer. In previous definitions of rank, the value typically depended only on the current structure of the heap (One exception to this has been in order to get better bounds on INSERT, nodes are treated differently for potential purposes depending on whether or not they will ever be deleted. See [18, 12, 11]). Here, however, the definition is more nuanced in that for the purposes of the analysis only, nodes are classified into *marked* and *unmarked* categories based on the history of the structure, and these marks, along with the current structure of the heap, are used to compute the rank of each node.

For ease of presentation, the rank of a node is defined in terms of the function $j(n) = 2dc(n)+1$.

The general idea is to have the rank of a node be maintained so it is the negation of the key value stored in the node. (Ranks will be non-negative, and we will only give nodes non-positive integer key values; these can be assumed to be perturbed arbitrarily to give a total ordering of key values). A node's rank can increase as the result of a pairing, and a node's value can decrease as the result of a DECREASE-KEY. It is thus our goal to perform a DECREASE-KEY on a node which has had its rank increase to restore it to the negation of its rank. During the time between when a rank increase occurs in a node and the time the DECREASE-KEY is performed, we refer to the node as *marked*.

Call the *unmarked subtree* of a node to be the subtree of a node if all marked nodes were detached from their parents; the *unmarked structure* of the heap is the structure of the unmarked subtrees of the roots. The rank of a node at a given time will be defined as a function of the structure of its unmarked subtree.

We emphasize that the notion of marking a node is for the purposes of the analysis only, such marks need not be stored.

The following assumes a particular heap structure and marking, as the rank of a node is always defined with respect to the structure of the heap after executing a sequence of heap operations.

Let x be the node we wish to compute the rank of. Let k denote the number of unmarked children of x, and let $y_1, y_2, \ldots y_k$ denote these children numbered right-to-left (i.e., in the order which they became children of x).

Let $\tau_i(x)$ be a subtree of x consisting of x connected to only the subtrees induced by $y_1, y_2, \ldots y_i$. We will define the function $r_i(x)$ as a function of $\tau_i(x)$. The rank of a node, $r(x)$ is $r_k(x)$.

Each node y_i is said to be *efficiently linked* to its parent x if and only if $r_{i-1}(x) - c \leq r(y_i) \leq r_{i-1}(x)$. The case of $r(y_i) > r_{i-1}(x)$ will never occur, as pairings will only happen among unmarked nodes, where the rank perfectly matches the negation of the key value.

We will have the property that $r_i(x)$ is either $r_{i-1}(x)$ or $r_{i-1}(x) + 1$; in the latter case, y_i is called *incremental*.

Given a node y_i, let j be defined to be the index of the first incremental node in the sequence $\langle y_{i-1}, y_{i-2}, \ldots \rangle$; j is defined to be 0 if there is no such incremental node. The set $N(y_i)$ is defined to be $\{y_k | j < k \leq i\}$; that is, y_i and the maximal set of its non-incremental siblings to the right.

Let c be a constant to be determined later in Section 3.6. Given these preliminaries, we can now give the full definition of the rank of a node:

$$r_i(x) = \begin{cases} 0 & \text{if } i = 0 \\ r_{i-1}(x) + 1 & \begin{array}{l}\text{(Efficient case) } y_i \text{ is efficiently linked and is the } j(n)\text{th efficient element of } N(y_i) \\ \text{—or—} \\ \text{(Default case) } |N(y_i)| = j(n)^c\end{array} \\ r_{i-1}(x) & \text{otherwise} \end{cases}$$

While the rank and mark are interrelated, there is no circularity in their definitions—whether a node is marked depends on its rank and key value and the rank of a node is a function of the ranks and marks of its children.

Observation 2. *Given two nodes x and y with different ranks, the unmarked structure of their induced subtrees must be different.*

This follows directly from the fact that the rank of a node is a function of its induced unmarked subtree.

Definition 3. *Set $f(n) = \frac{j(n)^c}{2 \log j(n)}$ and $g(n) = \frac{j(n)}{2 \log j(n)}$.*

These two functions are defined for convenience. Note that $f(n) = o(\log n)$ easily since $dc(n) = o(\log \log n)$ and c is a constant.

Next, we present a key lemma that related the number of unmarked subtree size to the number of efficiently linked children.

Lemma 4. *Suppose a root with unmarked subtree of size m has $\leq f(n) \log m$ unmarked children. Then it has $\geq g(n) \log m$ efficiently linked children.*

Proof. We first prove that the size of an unmarked subtree rooted at a node of rank k is at most $j(n)^k$. Let s_k be the size of the maximum unmarked heap of rank k. Such a heap can be created

from a maximal unmarked heap of rank $k-1$, which has been paired to $j(n)-1$ maximal unmarked heaps of rank $k-1$ and $j(n)^c-(j(n)-1)$ maximal unmarked heaps of rank $k-c-1$. Thus:

$$s_k \leq (j(n)-1)s_{k-1} + (j(n)^c - (j(n)-1))s_{k-c-1}.$$

By induction,

$$s_k \leq (j(n)-1)j(n)^{k-1} + (j(n)^c - j(n)+1)j(n)^{k-c-1} \leq j(n)^k.$$

This implies that if there are m nodes in a node x's unmarked induced subtree, the rank of x is at least $\log_{j(n)} m$.

Suppose node v has rank $\geq k$ and at most $\frac{k}{2}j(n)^c$ unmarked children. Then, v has at least $kj(n)/2$ efficiently linked unmarked children, each having rank $< k$. To see this consider how the rank of a node is computed. If node v has less than $kj(n)/2$ efficiently linked children, then the rank increases due to the efficient case must be less than $k/2$ by the definition of rank. This implies that at least $k/2$ rank increases will be due to the default case. This implies that there are at least $\frac{k}{2}j(n)^c$ unmarked children, a contradiction.

Consider a root with unmarked subtree size m and thus rank at least $\log_{j(n)} m$. Suppose it has $\leq \frac{\log_{j(n)} m}{2} j(n)^c = f(n)\log m$ unmarked children. Then it has $\geq \frac{1}{2}j(n)\log_{j(n)} m = g(n)\log m$ efficiently linked unmarked children by the argument in the previous paragraph. □

3.1.2 Monotonic operation sequences

Call the *designated minimum root* the next node to be removed in an Extract-Min.

Definition 5 (Monotonic operation sequence)**.** *Define a* monotonic operation sequence *to be one where* Decrease-Key *operations are only performed on roots, children of the designated minimum root, or marked nodes.*

All of the sequences of operations we define will be monotonic. Recall the invariant that the rank of a node is the negation of its key value. The following lemma tells us that when the rank of a node changes, in order to recover the invariant, we can call Decrease-Key with a positive parameter Δx. The only case when this is not feasible is when the node is the designated minimum root. However, in this case the node will be removed by an extract-min operation.

Lemma 6 (Monotonic sequences and rank)**.** *In a monotonic operation sequence, the rank of a node never decreases, from the time it is inserted until the time it becomes the designated minimum root.*

Proof. In a monotonic operation sequence, for any node x with descendent y (at the beginning of an operation) where all nodes on the path from x's child down to and including y are unmarked, y will remain in the same location in x's subtree (at the beginning of subsequent operations) until x becomes the designated minimum root. Thus, since the rank of a node is a function of the structure of the unmarked subtree of the node which does not change in a monotonic operation sequence until the node becomes designated minimum root, the rank of a node never decreases until it becomes the designated minimum root. □

3.1.3 Augmented suboperations

We augment suboperation 6, the PAIR$(\cdot)$ operation, to return whether or not the rank was incremented as a result of the pairing. This augmentation does not give any more power to the pure heap model, since the exact ranks of all nodes is a function of the suboperation sequence and is thus known to the algorithm already. We use this augmentation to create a finer notion of what constitutes a distinct sequence of suboperations. In particular, we will use the following fact:

Lemma 7. *Suppose s_i and s_j are two structurally distinct states of the data structure. Suppose a single valid sequence of suboperations implementing an* EXTRACT-MIN *is performed on both, and the outcomes of all augmented suboperations that have return values are identical in both structures. Then, the position of all nodes who have had their ranks changed is identical in both.*

Proof. Observe that as a valid sequence of suboperations implementing EXTRACT-MIN is performed on both s_i and s_j, the labels of the pointers they perform each suboperation on are determined by a right to left order as defined in Section 2 and are thus equivalent. In addition, the only time a node can have its rank change is when something is paired to it, and this is now explicitly part of the return value of suboperation 6. □

3.1.4 Distinctness

Let $B = \langle b_1, b_2, \ldots \rangle$ be a sequence of priority queue operations. Let $A(b_i) = \langle a_1^i, a_2^i, \ldots \rangle$ be the sequence of augmented suboperations and their return values used by algorithm $\mathcal{A}$ to execute operation b_i if b_i is an EXTRACT-MIN; if it is not $A(b_i)$ is defined to be the empty sequence. $A(B)$ is the concatenation of $A(b_1), A(b_2), \ldots$.

We call two sequences of priority queue operations B and B' *algorithmically indistinguishable* if $A(B) = A(B')$, else they are *algorithmically distinct*.

Let $s_B(i)$ be the structure of the heap after running sequence $\langle b_1 \ldots b_i \rangle$; the *terminal structure* of B is $s_B(|B|)$ which we denote as s_B. Recall that by structure, we mean the raw shape of the heap without regard to the data in each node, but including which nodes are marked. Two sequences B and B' are *terminal-structure indistinguishable* if $s_B = s_{B'}$, else they are *terminal-structure distinct*. Given a set of mutually algorithmically indistinguishable and terminal-structure distinct (AI-TSD) sequences of heap operations Ξ, the *distinctness* of the set, $\xi(\Xi)$ is defined to be $\log |\Xi|$.

Note that having two sequences which are algorithmically indistinguishable does not imply anything about them being terminal-structure indistinguishable. For example, it may be possible to add a DECREASE-KEY to a sequence, changing the terminal structure, while the sequence of suboperations performed to execute the sequence remains unchanged. (Recall that suboperations only occur during EXTRACT-MIN operations).

A critical observation needed at the end of the proof is that the number of terminal-structure distinct sequences is function of n:

Lemma 8. *The maximum distinctness of any set Ξ of terminal-structure distinct sequences, all of which have terminal structures of size n, is $\xi(\Xi) = \mathcal{O}(n)$.*

Proof. The number of different shapes of a rooted ordered forest with n nodes is C_n, the n^{th} Catalan number. The number of different ways to mark some nodes in a forest with n nodes is 2^n. Since $C_n \leq 4^n$, the maximum number of distinct structures is at most $\log(2^n 4^n) = \mathcal{O}(n)$. □

3.2 Outline of the proof

In order to prove the result, we will show that assuming $dc(n) = o(\log \log n)$ leads to a contradiction, namely that of Lemma 8. We obtain this contradiction using the help of evolutions which will be described in detail in Section 3.3.

Evolutions take as input AI-TSD sets of sequences of heap operations and modify them by appending operations to all sequences or by removing some sequences. Evolutions are performed in rounds with the goal increasing the distinctness of the set of sequences of heap operations until a contradiction to Lemma 8 is obtained.

The only evolution that increases distinctness is the permutation evolution (Section 3.3.6) which given an AI-TSD set of sequences selects a set of nodes in the terminal structure of each sequence (let us call this set the permutable set), and appends sequences of Decrease-Key operations applied to every permutation of these nodes to increase the size of the set. Translating this increase in the size of the set to a contradiction of Lemma 8 comes with a few challenges.

First, the resulting set must be terminal-structure distinct. This is achieved by ensuring the permutable set consists of nodes with unique ranks, which by Observation 2 guarantees that all the new sequences generated are mutually terminal-structure distinct.

Second, the permutable set must be commonly located in the terminal structures of all sequences in the set. Thus, permutation evolution selects such a commonly located permutable set that is found in the terminal structures of a subset of the sequences, and removes the other sequences. This causes a decrease in distinctness.

Third, in order to have a non-negligible effect on distinctness, the permutable set must be large enough. This is achieved by considering the number of roots, which depends on how the algorithm executed Extract-Min during the previous round. We classify the terminal structures of sequences based on the number of the roots, the position of the root with the largest subtree, as well as the number of its children (Section 3.3.5). We pick the most common such root in all of the sequences and remove the rest of the sequences. This causes a decrease in distinctness. If the number of roots is small, this yields an adequate lower bound on the number of nodes in the largest subtree, and we perform permutation evolution. Otherwise, if the number of roots is large, we instead remove the root of the largest subtree by first performing a Decrease-Key on that root so that it becomes the designated minimum root, and then calling Extract-Min (Section 3.3.3). This increases the number of roots further before an Extract-Min is called (Section 3.3.4). Intuitively, since Extract-Min has to at least pair all the roots and since we assumed an amortized upper bound of $\alpha \log n$ on Extract-Min, we show that we cannot have a large number of roots in more than half of the rounds.

Lastly, calling Extract-Min causes a number of issues. First, the algorithm can execute a different sequence of suboperations (each of different lengths) to implement Extract-Min on each sequence. We select the most common sequence of suboperations and remove all others. This ensures that the resulting set of sequences are still mutually algorithmically indistinguishable. It also causes a decrease in distinctness. In addition, Extract-Min can cause rank increases leading to marked nodes. By Lemma 7, the positions of these nodes in the terminal structures of all of the sequences are identical. Therefore, since by Lemma 6 ranks never decrease, we can restore the invariant that rank of a node equals the negation of its key value by calling Decrease-Key on each such node.

At the end, we show that the decreases in distinctness due to removing sequences are dominated by the increase in distinctness that occurs during the permutation evolution. We also show that

the permutation evolution is invoked in at least half of the rounds (Lemma 10), implying that after sufficiently many rounds, the distinctness becomes $\omega(n)$, contradicting Lemma 8, and thus showing that the only assumption we used in the proof, $dc(n) = o(\log \log n)$, is false.

Next, we present the evolutions in detail in Section 3.3 and show how to group them into rounds in Section 3.4. We analyze the running time of executing a sequence formed after k rounds of evolutions in Section 3.5. Finally, we put it all together and prove the contradiction in Sections 3.6.

3.3 Evolutions

We will now describe several functions on AI-TSD sets of heap operations; we call such functions *evolutions*. The general idea is to append individual heap operations or small sequences of heap operations to all sequences in the input set Ξ and remove some of the resulting sequences so as to maintain the property that the sequences in the resultant set of sequences Ξ' are AI-TSD. The evolutions will also have the property that if the time to execute all sequences in Ξ is identical, then the runtime to execute all sequences in Ξ' will also be identical. The difference in the runtime to execute sequences in Ξ' versus those in Ξ will be called the *runtime* of an evolution.

3.3.1 Insert evolution

The *insert evolution* has the following form: $\Xi' = \text{EVOLVE-INSERT}(\Xi)$.

In an insert evolution, a single INSERT operation of a key with value 0 is appended to the end of all Ξ to obtain Ξ'. Given Ξ is AI-TDS, the set Ξ' is AI-TDS and trivially $\xi(\Xi) = \xi(\Xi')$. The runtime of the evolution is 1 since the added INSERT has runtime 1. The rank of the newly inserted node is 0, and is thus unmarked.

3.3.2 Decrease-key evolution

The *decrease-key evolution* has the following form: $\Xi' = \text{EVOLVE-DECREASE-KEY}(\Xi, p)$, where p is a location which is either a root or a marked node in all terminal structures of sequences in Ξ.

In a decrease key evolution, a $\text{DECREASE-KEY}(p, \Delta x)$ operation is appended to the end of all sequences in Ξ to obtain Ξ'. The value of Δx is chosen such that the new key value of what p points to is set to the negation of its current rank; this means Δx is always nonnegative because of the monotone property of ranks noted in Lemma 6. Observe that if p points to a marked node, then it is unmarked after performing an EVOLVE-DECREASE-KEY. This requirement ensures that all structures that are distinct before this operation will remain distinct after the operation. Thus, the set Ξ' is AI-TDS and trivially $\xi(\Xi) = \xi(\Xi')$. The runtime of the evolution is 1 since the added DECREASE-KEY has runtime 1.

3.3.3 Designated minimum root evolution

Letting r be the position of one root which exists in all terminal structures of Ξ, the *designated minimum root evolution* has the form $\Xi' = \text{EVOLVE-DESIGNATED-MINIMUM-ROOT}(\Xi, r)$.

In a designated minimum root evolution, a DECREASE-KEY operation on r to a value of negative infinity is appended to all sequences in Ξ to give Ξ'. It will always be the case that the (next) evolution performed on Ξ' will be an extract-min evolution; the root r, which is known as the *designated minimum root*, will be removed from all terminal structures of Ξ' in this subsequent

EVOLVE-EXTRACT-MIN. There is no change in distinctness caused by this operation: $\xi(\Xi) = \xi(\Xi')$. The runtime of the evolution is 1 since the added DECREASE-KEY has runtime 1.

3.3.4 Extract-min evolution

The *extract-min evolution* has the form $(\Xi', V, e) = \text{EVOLVE-EXTRACT-MIN}(\Xi)$.

The extract-min evolution is more complex than those evolutions previously described, and the derivation of Ξ' from Ξ is done in several steps.

First, an EXTRACT-MIN operation is appended to the end of all sequences in Ξ to obtain an intermediate set of sequences which we call $\overline{\Xi}$. Recall that the EXTRACT-MIN operation is implemented by a number of suboperations. There is no reason to assume that the suboperations executed by the algorithm in response to the EXTRACT-MIN in each of the elements of $\overline{\Xi}$ are the same; thus the set $\overline{\Xi}$ may no longer be algorithmically indistinguishable. We fix this by removing selected sequences from the set $\overline{\Xi}$ so that the only ones that remain execute the appended EXTRACT-MIN by using identical sequences of suboperations. This is done by looking at the first suboperation executed in implementation of EXTRACT-MIN in each element of $\overline{\Xi}$, seeing which suboperation is the most common, and removing all those sequences $\overline{\Xi}$ that do not use the most common first suboperation. If the suboperation is one which has a return value, the return value which is most common is selected and the remaining sequences are removed. This process is repeated for the second suboperation, etc., until the most common operation is END() and thus the end of all remaining suboperation sequences has been simultaneously reached. Since there are only a constant η number of suboperations, and return values, if present, are boolean, at most a constant fraction (specifically $1 - \frac{1}{2\eta}$) of $\overline{\Xi}$ is removed while pruning each suboperation. At the end of processing each suboperation by pruning the number of sequences, the new set is returned as Ξ'. The set Ξ' can be seen to be terminal-structure distinct, since pairing identically positioned roots in structurally different heaps, and having the same nodes win the pairings, can not make different structures the same.

Observe that the nodes winning pairings in the execution of the EXTRACT-MIN might have their ranks increase, and thus become marked. Moreover, due to Lemma 7, the position of all such nodes is identical in all terminal structures of Ξ'. The set of the locations of these newly marked nodes is returned as V, the *violation set*.

Observe that it has been ensured that all of the sets of operations execute the appended EXTRACT-MIN using the same suboperations. Thus, we define e to be this common number of suboperations used to implement the EXTRACT-MIN; this value is returned by the evolution. As each suboperation reduces the distinctness by at most a constant, $\xi(\Xi') \geq \xi(\Xi) - e\log(2\eta) = \xi(\Xi) - \mathcal{O}(e)$. The runtime of the evolution is e since that is the cost of the added EXTRACT-MIN.

3.3.5 Big/small evolution

The *big/small evolution* has the form $(\Xi', (p, \textit{bigsmall})) = \text{EVOLVE-BIG-SMALL}(\Xi)$.

The goal of the big/small evolution is to ensure that the terminal structures of all sets are able to be executed in the same way in subsequent evolutions. In a big/small evolution, each terminal structure of each of the operation sequences of Ξ is classified according to the following, using the previously-defined function $f(n)$:

- The exact number of roots (if at most $f(n)\log n$) or the fact that the number of roots is greater than $f(n)\log n$ (we call this case *many-roots*).

- If the exact number of roots is at most $f(n)\log n$:
  - The position of the root with the largest subtree (the leftmost such root if there is a tie). Call it p. Observe that the size of p's subtree is at least $\frac{n}{f(n)\log n}$.
  - The exact number of children of p if less than $f(n)\log\frac{n}{f(n)\log n}$ (we call this case *small*) or the fact that the number of roots is greater than $f(n)\log\frac{n}{f(n)\log n}$ (we call this case *root-with-many-children*).

There are at most $\lceil f(n)\log n\rceil^2\cdot\lceil f(n)\log\frac{n}{f(n)\log n}\rceil$ possible classifications. We create set Ξ' by removing from Ξ sequences with all but the most common classification of their terminal structures. The return value is based on the resultant classification:

Many-roots: Return $(p, \textit{bigsmall})$ where $p = \textsc{Null}$ and $\textit{bigsmall} = \textsc{Big}$.

Root-with-many-children: Return $(p, \textit{bigsmall})$ where p is the location of the root with the largest subtree and $\textit{bigsmall} = \textsc{Big}$.

Small: Return $(p, \textit{bigsmall})$ where p is the location of the root with the largest subtree and $\textit{bigsmall} = \textsc{Small}$.

We bound the loss of distinctness, which is the logarithm of the number of classifications. Since $f(n) = o(\log n)$, then $\log\left(\lceil f(n)\log n\rceil^2\cdot\lceil f(n)\log\frac{n}{f(n)\log n}\rceil\right) = \mathcal{O}(\log\log n)$, and thus $\xi(\Xi') = \xi(\Xi) - \mathcal{O}(\log\log n)$. The evolution's runtime is 0 since no operations are added to any sequence.

3.3.6 Permutation evolution

The *permutation evolution* has the form $\Xi' = \textsc{Evolve-Permute}(\Xi)$, where the leftmost root r has in all terminal structures of the sequences of Ξ a subtree size of at least $\frac{n}{f(n)\log n}$ and at most $f(n)\log\frac{n}{f(n)\log n}$ children; this will be achieved by being in the small case of the big/small evolution and performing a decrease-key evolution on the relevant node. It is required that all terminal structures of sequences in Ξ are entirely unmarked.

In a permutation evolution, distinctness increases, and is the only evolution to increase the number of sequences in the process of converting Ξ to Ξ'.

Combining the preconditions of the permutation evolution with Lemma 4, yields the fact that all nodes in the terminal structures of Ξ at location r have at least $g(n)\log\frac{n}{f(n)\log n}$ efficiently linked children; since there are at most $(c+1)j(n)$ efficiently linked children of each rank (due to the definition of efficiently linked), that means there are at least $\frac{g(n)\log\frac{n}{f(n)\log n}}{(c+1)j(n)}$ efficiently linked children of different ranks in each terminal structure. Find such a set and call it the *permutable set* (chose one arbitrarily if there is more than one possibility). Look at the position of these permutable sets in all terminal structures of the sequences of Ξ, and pick the position of the permutable set that is most common. Form the intermediate set of sequences $\hat{\Xi}$ by removing from Ξ all sequences that do not have this commonly located permutable set. Letting $F = f(n)\log\frac{n}{f(n)\log n}$ and $G = \frac{g(n)}{(c+1)j(n)}\log\frac{n}{f(n)\log n}$, an upper bound on the logarithm of the number of different locations a permutable set of size G could be in is

$$\log\binom{F}{G} = \Theta(G\log\frac{F}{G}) = \Theta\left(\frac{\log n}{\log j(n)}\cdot\log j(n)\right) = \Theta\left(\log n\right)$$

given $f(n) = o(\log n)$ and $g(n) = \Theta(j(n)/\log j(n))$. Thus this step decreases the distinctness of the set by at most the logarithm of the number of commonly located permutable sets:

$$\xi(\hat{\Xi}) - \xi(\Xi) = -\mathcal{O}(\log n).$$

The permutable set is of size $\frac{g(n)\log\frac{n}{f(n)\log n}}{(c+1)j(n)}$. Using the definitions of $f(n)$, $g(n)$, c, and $dc(n)$, this is $\Theta\left(\frac{\log n}{\log dc(n)}\right)$. Let m be a constant such that the permutable set is of size at least $\frac{m\log n}{\log dc(n)}$ for sufficiently large n. Recall that in all the terminal structures of the sequences of $\hat{\Xi}$, the ranks of the items in the permutable sets are different, and in fact are strictly increasing, when viewed right-to-left as children of r.

We then create Ξ' by replacing each sequence in $\hat{\Xi}$ with $\left(\frac{m\log n}{\log dc(n)}\right)!$ new sequences created by appending onto the end of each existing sequence a sequence of all possible permutations of DECREASE-KEY operations on all elements of an arbitrary subset of size $\frac{m\log n}{\log dc(n)}$ of the permutable set.

The fact that all of the sequences in $\hat{\Xi}$ have the same permutable sets ensures that all terminal structures in Ξ' are terminal-structure distinct. (Recall that Observation 2 says that different ranks imply different structures of induced subtrees). Thus, in this step distinctness increases by $\xi(\Xi') - \xi(\hat{\Xi}) = \log\left(\frac{m\log n}{\log dc(n)}\right)! = \Theta\left(\frac{\log n\log\log n}{\log dc(n)}\right)$.

Thus, combining all the steps of the permutation evolution bounds the total increase of distinctness by

$$\begin{aligned}
\xi(\Xi') - \xi(\Xi) &= \xi(\Xi') - \xi(\hat{\Xi}) + \xi(\hat{\Xi}) - \xi(\Xi) \\
&= \Theta\left(\frac{\log n\log\log n}{\log dc(n)}\right) - \mathcal{O}(\log n) \\
&= \Theta\left(\frac{\log n\log\log n}{\log dc(n)}\right) && \text{since } dc(n) = o(\log\log n)
\end{aligned}$$

The cost of the evolution is at most $\frac{m\log n}{\log dc(n)}$, since that is the number of DECREASE-KEY operations appended to the sequences, and these all have unit cost.

3.4 Rounds

We proceed to perform a sequence of evolutions $\Psi = \langle\psi_0, \psi_1, \ldots\rangle$ to define a sequence of AI-TSD sets $\langle\Xi_0, \Xi_1, \ldots\rangle$. The initial set Ξ_0 consists of a single sequence of operations: the operation INSERT(0), executed n times. Each subsequent AI-TSD set Ξ_i is derived from Ξ_{i-1} by performing the single evolution ψ_{i-1}; thus in general Ξ_i is composed of some of the sequences of Ξ_{i-1} with some operations appended.

These evolutions are split into *rounds*; $\circ_i$ is the index of the first AI-TSD set of the ith round. Thus round i begins with AI-TSD set $\Xi_{\circ_i}$ and ends with $\Xi_{\circ_{i+1}-1}$ through the use of evolutions $\langle\psi_{\circ_i}\ldots\psi_{\circ_{i+1}-1}\rangle$

These rounds are constructed to maintain several invariants:

- All terminal structures of all sequences in the AI-TSD set at the beginning and end of each round have size n. This holds as in each round, exactly one insert evolution and exactly one extract-min evolution is performed.

Algorithm 1 Algorithmic presentation of how evolutions are used to build the sequence of AI-TSD sequences $\langle \Xi_0, \Xi_1, \Xi_2, \ldots \rangle$, which are split into rounds where the index of the start of round i is $\circ_i$.

$\Xi_0 = \{\overbrace{\langle \textsc{Insert}(0), \textsc{Insert}(0), \ldots, \textsc{Insert}(0) \rangle}^{n \text{ INSERT operations}}\}$
$i = 0$
$round = 0$
$\circ_0 = 0$;
loop
  $(\Xi_i, p, bigsmall) = \textsc{Evolve-Big-Small}(\Xi_{i++})$;
  if $bigsmall = \textsc{Small}$ **then**
    $\Xi_i = \textsc{Evolve-Decrease-Key}(\Xi_{i++}, p)$; ▷ Small round
    $\Xi_i = \textsc{Evolve-Permute}(\Xi_{i++})$;
  else
    $\Xi_i = \textsc{Evolve-Designated-Minimum-Root}(\Xi_{i++}, p)$; ▷ Big round
  end if
  $(\Xi_i, V, e) = \textsc{Evolve-Extract-Min}(\Xi_{i++})$; ▷ Common to small and big rounds
  for each v in V **do**
    $\Xi_i = \textsc{Evolve-Decrease-Key}(\Xi_{i++}, v)$;
  end for
  $\Xi_i = \textsc{Evolve-Insert}(\Xi_{i++})$;
  $\circ_{++round} = i$;
end loop

- All nodes in all terminal structures in the AI-TSD sets at the beginning and end of each round are unmarked.

There are two types of rounds, *big rounds* and *small rounds*. At the beginning of both types of rounds a big/small evolution is performed. The return value of the big/small evolution determines whether it will be a big or a small round.

The reader may refer to Algorithm 1 for a concise presentation of how evolutions are used to construct $\langle \Xi_0, \Xi_1, \Xi_2, \Xi_3, \ldots \rangle$. We now describe this process in detail.

3.4.1 The Big Round

As the round begins, we know that the terminal structures of the AI-TSD set are entirely unmarked, and there are either at least $f(n) \log n$ roots, or one root with at least $f(n) \log \frac{n}{f(n) \log n}$ children. The round proceeds as follows:

1. Perform a designated minimum root evolution on the root with largest subtree; this is the node r from the big/small evolution whose location is encoded in the return value; as a result of the big/small evolution it is guaranteed to be in the same location in all of the terminal structures of the sequences of Ξ.

2. Perform an extract-min evolution.

3. For each item in the violation sequence returned by the extract-min evolution, perform a decrease-key evolution. This makes the terminal structures of all heaps in Ξ unmarked.

4. Perform an insert evolution.

Assuming we are in round i, let e_i be the cost of the extract-min evolution, and let v_i be the size of the violation sequence. The cost of the round (the sum of the costs of the evolutions) is e_i+v_i+2, which is at least $f(n)\log\frac{n}{f(n)\log n}$, and based on the evolutions performed the distinctness can be bounded as follows:

$$\xi_i - \xi_{i+1} = \overbrace{\mathcal{O}(e_i)}^{\textsc{Evolve-Extract-Min}} + \overbrace{\mathcal{O}(\log\log n)}^{\textsc{Evolve-Big-Small}}.$$

3.4.2 The Small Round

In a small round, there is one root, call it x, at the same location in all terminal structures, with size at least $\frac{n}{f(n)\log n}$ and some identical number of children in all terminal structures which is at most $f(n)\log\frac{n}{f(n)\log n}$. The location of x was returned by the big/small evolution. The round proceeds as follows:

1. Perform a decrease-key evolution on x to make it negative infinity.
2. Perform an permutation evolution.
3. Perform an extract-min evolution.
4. For each item in the violation sequence returned by the extract-min evolution, perform a decrease-key evolution.
5. Perform an insert evolution.

Let e_i be the actual cost of the Extract-Min, let v_i be the size of the violation sequence. The cost of the round is $e_i + v_i + 2 + \frac{m\log n}{\log dc(n)}$, and based on the evolutions performed the distinctness can be bounded as follows:

$$\xi_i - \xi_{i+1} = \overbrace{\mathcal{O}(e_i)}^{\textsc{Evolve-Extract-Min}} + \overbrace{\mathcal{O}(\log\log n)}^{\textsc{Evolve-Big-Small}} - \overbrace{\Omega\left(\frac{\log n\log\log n}{\log dc(n)}\right)}^{\textsc{Evolve-Permute}}.$$

3.5 Upper Bound on Time

The following lemma is needed in the next section.

Lemma 9. *The total time to execute any sequence in $\Xi_{\circ_k}$ is $\mathcal{O}\left(\frac{dc(n)}{\log dc(n)}\cdot k\log n\right)$.*

Proof. Let B be a sequence in $\Xi_{\circ_k}$, and let $R(B)$ be the time to execute B. Let dc_i be the cost of the permutation evolution in round i; this is at most $\frac{m\log n}{\log dc(n)}$ if round i is a small round and 0 if round i is a big round (recall that permutation evolutions are only performed in the small round). Thus the cost for any round i, whether big (§3.4.1) or small (§3.4.2), can be expressed as $2+dc_i+e_i+v_i$. The time to execute any sequence $B\in\Xi_{\circ_k}$, which we denote as $R(B)$, is the sum of the costs of the rounds:

$$R(B) = \sum_{i=1}^{k}(2 + dc_i + e_i + v_i) \tag{1}$$

An item in the violation sequence has had its rank increase. Its rank can only increase after winning $j(n)$ pairings. Pairings are operations. Thus,

$$\sum_{i=1}^{k} v_i \leq \frac{R(B)}{j(n)} \tag{2}$$

Combining (1) and (2) and rearranging gives:

$$\begin{aligned}
\sum_{i=1}^{k} v_i &\leq \frac{1}{j(n)}\sum_{i=1}^{k}(2 + dc_i + e_i + v_i) \\
\left(1 - \frac{1}{j(n)}\right)\sum_{i=1}^{k} v_i &\leq \frac{1}{j(n)}\sum_{i=1}^{k}(2 + dc_i + e_i) \\
\sum_{i=1}^{k} v_i &\leq \frac{1}{j(n)-1}\sum_{i=1}^{k}(2 + dc_i + e_i) \qquad (3)
\end{aligned}$$

We know by assumption that the runtime is bounded by the sum of the amortized costs stated at the beginning of §3. Each round has one INSERT and one EXTRACT-MIN (at an amortized cost of $\alpha \log n$ each) and $1 + dc_i + v_i$ DECREASE-KEY operations (at an amortized cost of $dc(n)$ each). This gives:

$$R(B) \leq \sum_{i=1}^{k}((dc_i + v_i + 1)dc(n) + 2\alpha \log n) \tag{4}$$

Combining (3) and (4) gives:

$$R(B) \leq \sum_{i=1}^{k}\left[\left(dc_i + \frac{2 + dc_i + e_i}{j(n)-1} + 1\right)dc(n) + 2\alpha \log n\right]$$

Since $dc_i \leq \frac{m \log n}{\log dc(n)}$ (recall that m is a constant defined in §3.3.6).

$$R(B) \leq \sum_{i=1}^{k}\left[\left(\frac{2}{j(n)-1} + 1\right)dc(n) + \frac{e_i dc(n)}{j(n)-1} + 2\alpha \log n + \left(1 + \frac{1}{(j(n)-1)}\right)\frac{mdc(n)}{\log dc(n)} \log n\right]$$

$$R(B) \leq k \log n\left[\left(\frac{2}{j(n)-1} + 1\right)\frac{dc(n)}{\log n} + 2\alpha + \left(1 + \frac{1}{j(n)-1}\right)\frac{mdc(n)}{\log dc(n)}\right] + \frac{dc(n)}{j(n)-1}\sum_{i=1}^{k} e_i$$

Since $\sum_{i=1}^{k} e_i \leq R(B)$

$$R(B)\left(1-\frac{dc(n)}{j(n)-1}\right) \le k\log n\left[\left(\frac{2}{j(n)-1}+1\right)\frac{dc(n)}{\log n}+2\alpha+\left(1+\frac{1}{j(n)-1}\right)\frac{mdc(n)}{\log dc(n)}\right]$$

Substituting in the definition of $j(n)$: $j(n) = 2dc(n)+1$

$$R(B)\left(1-\frac{1}{2}\right) \le k\log n\left[\left(\frac{1}{dc(n)}+1\right)\frac{dc(n)}{\log n}+2\alpha+\left(1+\frac{1}{2dc(n)}\right)\frac{mdc(n)}{\log dc(n)}\right]$$

$$R(B) \le 2k\log n\left[\frac{1}{\log n}+\frac{dc(n)}{\log n}+2\alpha+\frac{mdc(n)}{\log dc(n)}+\frac{m}{2\log dc(n)}\right]$$

Since $dc(n) = o(\log\log n)$, for large enough n,

$$R(B) \le 2k\log n\left[2\alpha+\frac{mdc(n)}{\log dc(n)}+\frac{m}{2\log dc(n)}\right] < 4k\log n\left[\alpha+\frac{mdc(n)}{\log dc(n)}\right]. \tag{5}$$

This yields $R(B) = \Theta\left(\frac{dc(n)}{\log dc(n)}\cdot k\log n\right)$.

□

3.6 Obtaining a contradiction

Lemma 10. *After k rounds, at least $\frac{k}{2}$ rounds must be small rounds.*

Proof. Proof is by contradiction. Suppose more than $\frac{k}{2}$ rounds are big rounds. The actual cost of a big round is at least $f(n)\log\frac{n}{f(n)\log n}$, so the total actual cost is greater than:

$$\begin{aligned} R(B) &\ge \frac{kf(n)}{2}\log\frac{n}{f(n)\log n} \\ &= k\frac{j(n)^c}{4\log j(n)}\log\frac{2n\log j(n)}{j(n)^c\log n} && \text{since } f(n)=\frac{j(n)^c}{2\log j(n)} \end{aligned}$$

since $j(n) = 2dc(n)+1$ and $dc(n) = o(\log\log n)$, for large enough n we have

$$\begin{aligned} &> k\log n^{2/3}\cdot\frac{2^c dc(n)}{4\log dc(n)^{4/3}} \\ &= k\log n\cdot\frac{2^c dc(n)}{8\log dc(n)} \end{aligned}$$

Setting $c = 1+\log(32\alpha+32m)$ yields

$$\begin{aligned} &> k\log n\left[\frac{32\alpha dc(n)}{8\log dc(n)}+\frac{32mdc(n)}{8\log dc(n)}\right] \\ R(B) &> 4k\log n\left[\alpha+\frac{mdc(n)}{\log dc(n)}\right] \end{aligned}$$

This contradicts Equation (5) in Lemma 9. Thus, at least $\frac{k}{2}$ rounds must be small rounds. □

Proof of Theorem 1. The distinctness gain of a round has been bounded as follows:

$$\xi_{\circ_{i+1}} - \xi_{\circ_i} = \begin{cases} -\mathcal{O}(e_i) - \mathcal{O}(\log\log n) & \text{If the } i\text{th round is a big round (§3.4.1)} \\ -\mathcal{O}(e_i) - \mathcal{O}(\log\log n) + \Theta\left(\frac{\log n \log\log n}{\log dc(n)}\right) & \text{If the } i\text{th round is a small round (§3.4.2)} \end{cases}$$

Now we know that $\sum_{i=1}^{k} e_i$ is less than the actual cost to execute a sequence in $\Xi_{\circ_k}$, which is $\mathcal{O}\left(\frac{dc(n)}{\log dc(n)} \cdot k \log n\right)$ by Lemma 9. Substituting $\sum_{i=1}^{k} e_i = \mathcal{O}\left(\frac{dc(n)}{\log dc(n)} \cdot k \log n\right)$ into the above and using the fact from Lemma 10 that at least half of the rounds are small rounds gives:

$$\xi_{\circ_k} - \xi_{\circ_0} = \Theta\left(\frac{\log\log n}{\log dc(n)} \cdot k \log n\right) - \mathcal{O}(k \log\log n) - \mathcal{O}\left(\frac{dc(n)}{\log dc(n)} \cdot k \log n\right) \tag{6}$$

Since $dc(n) = o(\log\log n)$, the negative terms in the previous equation can be absorbed, giving:

$$\xi_{\circ_k} - \xi_{\circ_0} = \Theta\left(\frac{k \log n \log\log n}{\log dc(n)}\right) = \omega\left(k \cdot \frac{\log n \log\log n}{\log\log\log n}\right).$$

But after sufficiently many rounds (i.e. sufficiently large k) this contradicts Lemma 8 that for all i, $\xi_i = \mathcal{O}(n)$. Thus for sufficiently large k and n a contradiction has been obtained, proving Theorem 1. □

References

[1] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms (3. ed.)*. MIT Press, 2009.

[2] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271, 1959.

[3] A. Elmasry. Pairing heaps with $O(\log\log n)$ decrease cost. In *SODA*, pages 471–476, 2009.

[4] A. Elmasry. Pairing Heaps with Costless Meld. In *ESA*, pages 183–193, 2010.

[5] A. Elmasry. The Violation Heap: a Relaxed Fibonacci-like Heap. *Discrete Math., Alg. and Appl.*, 2(4):493–504, 2010.

[6] M. L. Fredman. A Priority Queue Transform. In *WAE*, pages 244–258, 1999.

[7] M. L. Fredman. On the Efficiency of Pairing Heaps and Related Data Structures. *J. ACM*, 46(4):473–501, 1999.

[8] M. L. Fredman, R. Sedgewick, D. D. Sleator, and R. E. Tarjan. The Pairing Heap: A New Form of Self-Adjusting Heap. *Algorithmica*, 1(1):111–129, 1986.

[9] M. L. Fredman and R. E. Tarjan. Fibonacci heaps and their uses in improved network optimization algorithms. *J. ACM*, 34(3):596–615, 1987.

[10] B. Haeupler, S. Sen, and R. E. Tarjan. Rank-Pairing Heaps. *SIAM J. Comput.*, 40(6):1463–1485, 2011.

[11] J. Iacono. Improved Upper Bounds for Pairing Heaps. In *SWAT*, pages 32–45, 2000.

[12] J. Iacono. Improved Upper Bounds for Pairing Heaps. *CoRR*, 2011.

[13] J. Iacono and Ö. Özkan. Why some heaps support constant-amortized-time decrease-key operations, and others do not. In *ICALP*, pages 637–649, 2014.

[14] H. Kaplan and R. E. Tarjan. Thin heaps, thick heaps. *ACM TALG*, 4(1), 2008.

[15] S. Pettie. Towards a Final Analysis of Pairing Heaps. In *FOCS*, pages 174–183, 2005.

[16] D. D. Sleator and R. E. Tarjan. Self-Adjusting Binary Search Trees. *J. ACM*, 32(3):652–686, 1985.

[17] D. D. Sleator and R. E. Tarjan. Self-Adjusting Heaps. *SIAM J. Comput.*, 15(1):52–69, 1986.

[18] J. T. Stasko and J. S. Vitter. Pairing Heaps: Experiments and Analysis. *Commun. ACM*, 30(3):234–249, 1987.

[19] R. E. Wilber. Lower Bounds for Accessing Binary Search Trees with Rotations. *SIAM J. Comput.*, 18(1):56–67, 1989.